Sponsored by
Solera Networks

When Breaches Happen: Top Five Questions to Prepare For

June 2012

A SANS Whitepaper

Written by: Dave Shackleford

Summary: Plan for the Inevitable

New technology, the web, mobility, social media and applications are driving business to new heights, but also exposing new risks. Data expansion/propagation is exploding, and the surface area for attackers is greater than ever before. Meanwhile, attackers are more sophisticated, their exploits are lingering longer, and IT security response is being stretched beyond current capabilities. According to the latest Verizon Data Breach report, 85 percent of organizations breached did not realize they had been compromised for weeks or longer, and they usually only found out after an external third party alerted them.[1] For example, Krebs on Security reported Global Payment Systems had been infiltrated for more than six months when, in February, it discovered the loss of 1.5 million cardholder records to thieves.[2] And, as reported by The Wall Street Journal, cybercriminals enjoyed unhindered access to the corporate network of Nortel Networks for nearly a decade, beginning as early as 2000.[3]

Why are organizations having so much trouble detecting incidents in their environments? The reasons are many, including:

- Adversaries continue to grow more sophisticated and can convincingly target individual users and specific systems while hiding their activities from sensors.
- Hacktivists and politically motivated groups, such as Anonymous and Lulzsec, now relentlessly search for and take advantage of multiple attack surfaces, and then make their attacks public.
- Organizations have too much event data to sort through and simply can't find events on their networks, according to this year's SANS Annual Log Management Survey.[4]
- Today's over-focus on preventative technologies--Intrusion Prevention Systems (IPSs), Next-Generation Firewalls (NGFWs), Data Leakage Prote ction (DLP), and malware detection—is critical; but, the technologies fail to see and block advanced malware and advanced targeted attacks due to shortcomings such as deployment constraints, configuration errors or "event blindness" (during the period when no signatures are available for a new attack).
- Security Information and Event Management (SIEM), while effective at presenting and correlating log and event data, can only deliver information that is as complete as what the SIEM receives. Reviewing SIEM data is like reviewing a phone bill and seeing who talked to whom and at what time, but not having the actual conversation. Not knowing the conversation, responders cannot determine impact, intent and other vital information needed when events occur.

To put it bluntly, most organizations ***will*** be breached at some point and should plan for this inevitability. To prepare, IT security and operations groups need to put technology and processes in place to help them quickly sort through all their event data to determine who was behind any security breach—and at the same time uncover what systems were compromised, what data was extricated, whether the attackers are still on the network or not, and how to make improvements to assure this exact breach method won't succeed again. This paper explores how to create these processes so that in the event of a breach, IT security and operations teams are ready to respond with actionable information.

---

1 www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf

2 http://krebsonsecurity.com/2012/05/global-payments-breach-window-expands/

3 http://on.wsj.com/Jr815y

4 www.sans.org/reading_room/analysts_program/SortingThruNoise.pdf

Prevention Ideal. Response Inevitable.

Breaches happen. Preparation, then, is a matter of where to put your efforts. Organizations need to prepare for the reality that major security events will happen—including breaches that can impact the risk and compliance posture of the organization.

By all means, security professionals should make every effort to prevent incidents, but it's not realistic to assume that such efforts can be universally effective and negate the need for established incident response processes. Responders need a plan of action going into the breach—and tools to support those actions in a timely, auditable fashion.

In November 2011, John Kindervag and Rick Holland of Forrester Research described the shortcomings in a prevention-only mindset in their survey-based paper, "Planning for Failure."[5] In it, Forrester surveys showed that 25 percent of respondents experienced a breach of some kind in the last 12 months, yet 30 percent of those breached made no changes at all to their security programs, and only 18 percent of breached organizations increased spending on incident response.

This is not a realistic approach given today's advanced threats, which are harder than ever for organizations to detect and respond to, according to more than 600 respondents who completed the SANS 8th Log and Event Management Survey (See Figure 1).[6]

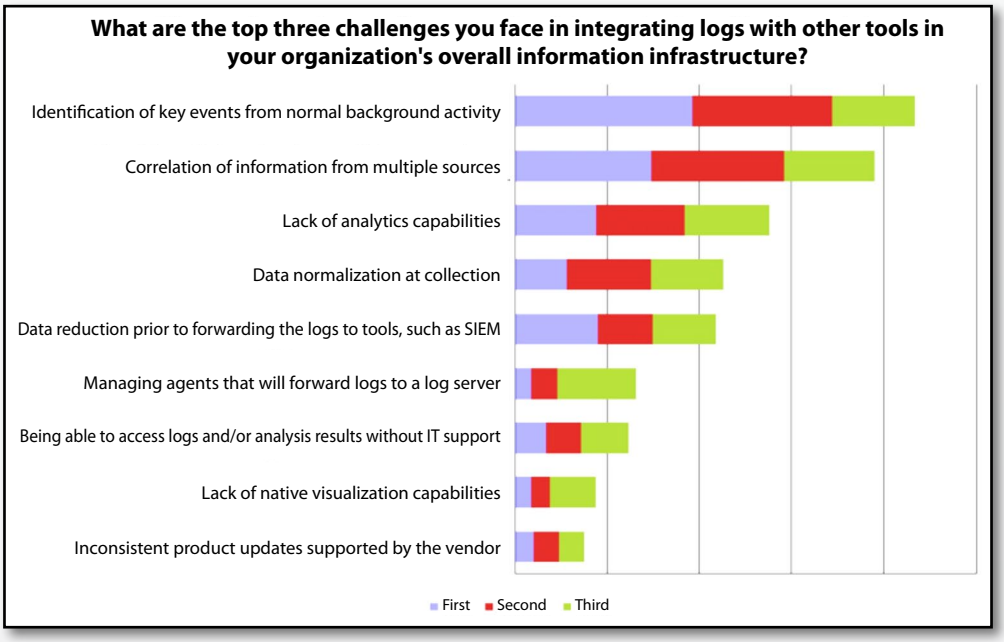

Figure 1. SANS Survey Responses: Event Detection and Correlation Difficulties

The reasons organizations have more difficulty finding events goes beyond the advanced nature of today's threat and into problems with too much data to sort and analyze, according to the SANS survey and other reports. However, given the appropriate analytics, this very data—or rather the context it provides—can be used to learn and respond to events faster and with more accuracy.

5 http://blogs.forrester.com/rick_holland/11-11-09-planning_for_failure

6 www.sans.org/reading_room/analysts_program/SortingThruNoise.pdf

The Big Picture

Determining where the breach occurred and how attackers entered the network usually comes down to researching several specific data sources, including the following:

- Log data for systems, applications, network devices and access control systems
- Active network connections
- Network traffic captures—including all relevant traffic, files, packets and metadata flowing into and out of the network
- Security event data from antivirus programs, Intrusion Detection Systems, firewall events, and so on)

Threat Analysis Challenges

Threat analysis can be performed with specialized tools, such as SIEM platforms and malware analysis sandbox tools. However, there are a number of challenges in today's complex environments that make such solutions less tenable for security incident detection and analysis:

- **The "Big Data" problem** – In March 2012, Gartner analyst Neil MacDonald published a report called "Information Security Is Becoming a Big Data Analytics Problem."[7] In this report, MacDonald notes that businesses have a staggering array of security data, ranging across network packet data, multisource security event data, monitoring information, account management logs and more. Traditional vendor products are having a hard time scaling to manage this level of volume, and organizations are not pinpointing the data they need, according to the SANS Log and Event Management Survey.

  MacDonald suggests that security analytics platforms will become more and more critical, especially in determining that traditional preventative controls have been bypassed and advanced attacks are underway. MacDonald's position is that tools that provide contextual security data with algorithm-driven analysis, and analysts who understand how to use these kinds of tools, will become standard elements of enterprise security programs in the near future.

- **More advanced threats** – Concurrently, threats are growing more advanced, making them more subtle and difficult to detect. The model of looking for "bad things" is becoming untenable because today's attacks blend in and hide so well. An attack may stay resident in memory or integrate with the OS kernel, making detection more difficult with traditional antimalware tools.

  A better bet for longer-term analysis is to build a baseline of known good behavior in the environment. This step adds situational context to the data and events to enable security teams to be more effective at spotting attacks as they're happening. However, capturing, inspecting, indexing and warehousing this volume of security events and big data is a tall order for organizations trying to integrate this data.

7 www.gartner.com/id=1960615

- **Less operational capacity in security** – Security teams often lack the time and manpower to adequately monitor numerous consoles with disparate data sets. Yet, as the SANS Log and Event Management Survey indicates, there is a driving need to bring together multiple data sources to analyze behavioral patterns and trends quickly and efficiently in order to detect attack indicators.
- **The diminishing effectiveness of predictive/preventative security controls** – Despite a robust security landscape, major security events and breaches continue at an increasingly more rapid pace. Although today's next-gen security products—including next-gen firewalls, Intrusion Prevention Systems (IPSs), SIEMs and malware detection platforms—are working to minimize gaps, they can't close them. Today's advanced malware and sophisticated attacks unfortunately fly under the radar of predictive, signature-based prevention tools. They simply can't stop what they can't see.

Visibility of Everything

A technique gaining a lot of traction in solving these problems is security intelligence and analytics (SIA). SIA focuses on scrutinizing large volumes of security data, emphasizing full packet capture, analysis, and correlation with logs and other event data.

The goal of security intelligence and analytics is to develop context around security events and information. This allows security teams to more thoroughly and effectively monitor the environment—while spotting unusual and potentially malicious events. This approach gives information security groups the ability to thoroughly investigate and understand "known events" that present clear indicators of compromise either through automated reporting (such as an IPS alert), or human detection (such as a call to a helpdesk). SIA also provides a mechanism to deliver risk-based security intelligence around "unknown events" that can hang out and evade detection for periods of time.

With the right context, big data security analytics can bring together the disparate pieces of an event to determine impact, scope, regulatory posture, and appropriate steps for remediation.

For example, root cause analysis can be used to determine whether breached systems or environments were compliant with specific controls, such as data encryption and retention, at the time of compromise. In addition, analytics can validate that firewall rules and other best practice controls were in place and functioning at the time of breaches (thus reducing the burden of compliance penalties on the victimized organizations). Audit trail data can also be validated or investigated during audits or prior to audits, allowing organizations to take a more proactive stance on meeting compliance controls and maintaining them over longer periods.

7 www.gartner.com/id=1960615

Focus on Preparedness: Top Five Questions to Answer After a Security Breach

IT professionals have at their disposal a vast array of security data sources that can provide the context needed to tell good from bad, normal from abnormal and other indicators of something illicit happening on the network. From the data derived from these sources, security intelligence and analytics can help them reconstruct events, files, traffic flows and even individual packets to answer the following five critical questions that must be asked in the event of a breach.

1. What Systems and Data Were Affected?

Determining system impact and data affected is the number one priority for organizations to both stop the bleeding and accurately report breaches to regulators within a reasonable timeframe. Generally, this involves exhaustive local system and application checks and forensic examination of network traffic and system/application data to determine what was accessed.

For many organizations, this can be a daunting task. In some cases, organizations may never fully determine the extent of a breach—due to data loss or corruption, increasingly automated host anti-forensic techniques deployed by attackers, or incomplete evidence and artifacts to support the investigation. If they're stealthy, attackers often have ample opportunity to clean up after themselves, remove logs, delete tools and cover their tracks at the system and application levels so that forensics cannot paint an accurate picture of events.

Using high-speed data analytics, response and audit professionals can reconstruct and revisit the time period during which the attack or data compromise occurred and determine the scope of the issue by assessing traffic into and out of the systems and applications in question. To start, NetFlow and other baseline data can also be used to gain some context to the initial compromise and determine whether other systems and data stores have been accessed. If this can be done quickly enough, they may prevent additional data from being accessed or exfiltrated.

2. How Did They Do It?

This is the most important question for improving post-breach security posture and processes. During a breach investigation, responders need concrete evidence of what happened, what controls were bypassed or obviated, and how and where the organization must improve its methods for a better security and compliance posture. This involves several steps:

- Identify evidence and artifacts of the breach itself, looking at stored data such as logs and packet captures, as well as "real-time" alerting measures like Intrusion Detection System (IDS) events, if they exist.
- Trace back activity from the point of discovery to the point of entry.
- Develop context around this security data to understand what is happening. For example, if two systems on an internal network are suddenly communicating a lot and using unusual traffic patterns and behaviors, does this mean that one or both of the systems are compromised? Or does it mean that a developer made a configuration change that is causing an application to behave unpredictably?

With context-based analytics, responders and auditors can return to the scene of the crime and determine what happened before, during and after the attack. For example, they may be able to trace back activity to a specific database server that had default credentials in place, and then further back to a database administrator's laptop that was missing patches or had evidence of social engineering activity. This would be the point of infection. Without the ability to trace back activity and reconstruct it, audit reports can never be completed accurately, and IT groups cannot repair the vulnerabilities that allowed the attack to succeed along its trajectory in the network. (For advice on analyzing known and unknown event data, read the "Analyzing Known and Unknown Data" section that appears after the five questions.)

3. Who Did This to Us?

Achieving attribution remains a seemingly ever-elusive holy grail of information security. However, as more advanced incidents succeed in breaching networks, it is important to determine at least the ***category*** of attacker:

- Is the event externally or internally based?
- What is the attacker's motivation? Is the victim organization being targeted specifically by criminal elements, by hacktivist groups or by a new form of malware or scenario that organizations need to plan for in their security practices?
- Where is the attacker located?

Knowing the basic answers to these questions determines how the organization investigates. For example if it is a Human Resource (HR) department issue that might turn into legal trouble, or if the case is likely to go to court, specific forensics practices must be included in the response. If it is an attack that occurs from outside the company, there may be limited recourse.

Determining who launched the attack may still be the most difficult question to answer definitively during or after a breach. Responders can return to specific IP addresses and account information found in logs and other security data and even to full packet captures to rebuild traffic sessions in the same time frame. They can also examine actual artifacts of specific activity. These activities allow for more effective backtracking from one event to another until the responders can determine the specific source of the attack and block future access.

4. Is It Really Over?

Effective monitoring for additional signs of malicious activity is key to determining whether the breach is fully remediated or if traces have been left behind. After the breach is identified, contained and repaired, continual monitoring for unusual events, account activity and network connections, as well as specific data patterns and packets in traffic, is key. Prematurely decommissioning systems is an error often committed in the course of incident response. While it seems the natural thing to do, this action clearly alerts attackers that they've been detected. They can, then, easily move to any number of additional points of presence hidden in the compromised network. This is why achieving a degree of certainty that the threat has been fully removed is so essential.

There are several important considerations when evaluating whether a breach is still underway, including:

- Be sure the cause of the breach was accurately identified. This is critical in determining the appropriate patterns of traffic and system behavior to look for.
- Have a prior understanding of what normal behavior is for both systems and network traffic. In order to determine this, teams need to start monitoring their environments much more diligently ***before*** they experience a breach.
- Once patterns of system and network activity have been identified, security teams can use analytics to compare pre- and post-breach behavior. This can lead to identification of other systems that need to be investigated, and security teams can follow up appropriately to make sure no additional malicious behavior is taking place.

5. Can It Happen Again?

Unfortunately, there is no foolproof solution for preventing breaches. Security teams can, however, develop a reasonable level of assurance that they've minimized specific exposure areas by doing a thorough analysis of their network environments before, during and after breach scenarios through the following steps:

- **Know and understand what is running in the environment and how it commonly behaves.** By developing sound baselines of network traffic and system and application behavior, organizations can more rapidly spot signs of malicious activity and respond more effectively.
- **Go over the root cause again.** Trace attacks back to their origin and ascertain what happened at the earliest possible time.
- **Reconstruct events.** Analytics that utilize full packet data (rather than just reading the headers) can help reconstruct events and deliver clear evidence.
- **Remediate security controls and improve user training.** Such improvements allow security and HR teams to reduce overall risk.

Analyzing Known and Unknown Data

The following two manufactured events provide examples of using analytics to leverage known and unknown event data to determine cause.

Analyzing Known Event Data

Security and network analytics will leverage "known" data, such as firewall and intrusion detection events. For example, the events in Figure 2 were generated from Cisco and NetScreen firewalls and show typical patterns that are indicative of normal, non-event activity.

```
(a) Aug 31 2011 07:45:33: %ASA-4-106023: Deny icmp src outside:5.6.7.8 dst
inside:1.2.3.4 (type 8, code 0) by access-group "outside_access_in"

(b) May 2 14:55:46 fire00 fire00: NetScreen device_id=fire00 [Root]system-
notification-00257(traffic): start_time="2012-05-02 14:55:45" duration=0
policy_id=119 service=udp/port:7001 proto=17 src zone=Trust dst
zone=Untrust action=Deny sent=0 rcvd=0 src=192.168.2.1 dst=1.2.3.4 src_
port=3036 dst_port=7001

(c) Sep 07 2011 09:12:25: %ASA-4-106023: Deny tcp src outside:1.2.3.4/4349
dst inside:5.6.7.8/139 by access-group "outside_access_in"
```

Figure 2: Normal Event Activity Derived from a Cisco ASA Firewall and a NetScreen Firewall

These events indicate (a) blocked access from an outside interface for ICMP traffic, (b) blocked access from an internal system trying to communicate outbound to a service on port 7001, and (c) a failed connection to Microsoft SMB services on port 139 within the environment. These represent traffic and system access patterns that can be monitored over time to develop a baseline of activity.

In addition, NetFlow records can determine which systems are communicating with one another. Network flow data can be used to develop a behavioral baseline of communications between systems in the same or different network segments (See Figure 3).

```
Sif  SrcIPaddress Dif  DstIPaddress     Pr SrcP DstP Pkts Octets StartTime
EndTime           Active  B/Pk  Ts  Fl

0059 127.0.0.1    005b 219.140.194.174  06 50   4f3   1   40
0721.21:58:00.593 0721.21:58:00.593 0.000   40    00  14

0059 127.0.0.1    005b 219.148.205.228  06 50   6ef   1   40
0721.21:57:56.533 0721.21:57:56.533 0.000   40    00  14
```

Figure 3: NetFlow Baseline Record of Normal Communications

As a final example of common known security data, network and host-based intrusion detection and prevention systems generate a plethora of information that can indicate known attack types or a variety of in-progress "events of interest." Figure 4 provides examples of simple IDS events.

```
(a) [**] [1:1807:10] WEB-MISC Chunked-Encoding transfer attempt [**]

[Classification: Web Application Attack] [Priority: 1]

05/02-13:20:05.804368 192.168.1.5:3602 -> 192.168.1.2:80

TCP TTL:128 TOS:0x0 ID:35630 IpLen:20 DgmLen:1500 DF

***A**** Seq: 0xBFF2387E  Ack: 0x9D37BACD  Win: 0xFAF0  TcpLen: 20

[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=2002-0392][Xref =>
http://cve.mitre.org/cgi-bin/cvename.cgi?name=2002-0079][Xref => http://
cve.mitre.org/cgi-bin/cvename.cgi?name=2002-0071][Xref => http://www.
securityfocus.com/bid/5033][Xref => http://www.securityfocus.com/bid/4485]
[Xref => http://www.securityfocus.com/bid/4474]


(b) [**] [1:621:7] SCAN FIN [**]

[Classification: Attempted Information Leak] [Priority: 2]

05/02-13:23:25.078708 192.168.1.57:57019 -> 192.168.1.2:723

TCP TTL:51 TOS:0x0 ID:25804 IpLen:20 DgmLen:40

*******F Seq: 0x7FBE0B31  Ack: 0x0  Win: 0x1000  TcpLen: 20

[Xref => http://www.whitehats.com/info/IDS27]

(c) 10.10.50.21 - - [02/May/2012:14:29:58 -0300] "GET /shell-cgi/ HTTP/1.1"
404 315 "-" "Mozilla/4.75 [en] (X11, U; Nessus)"
```

Figure 4: Known Host and Network Events Captured in IDS

In Figure 4, example (a) is a classic Apache web server attack signature from the Snort IDS, (b) is another Snort signature for a network scan, and (c) is an event from the host-based IDS OSSEC that indicates an attempt to access non-existent web server content.

All of these event types can be gathered and correlated in a network environment. Most mature security organizations are collecting and leveraging this data to some extent today. Additionally, while most organizations are collecting and leveraging this data, there is also a shift to full capture and collection of all network traffic to provide intelligence, analytics and full-fidelity reconstruction of events from these traditional sources.

Analyzing Unknown Events

Security teams face more difficulties in identifying "unknown" events, meaning those that don't have obvious signatures or identifiable patterns indicating malicious behavior. Today, network and security analytics can be performed with full packet capture data that allows longer-term behavioral profiling of the environment, as well as correlation with the "known" data discussed earlier. There are a number of benefits that full packet capture and analysis tools bring to the table for comprehensive incident detection and analysis, including:

- **Event Reconstruction** – If a specific behavioral indicator occurs in a certain time frame, security teams can reconstruct related traffic flows to determine what occurred while watching them after the fact in real-time.
- **File Extraction** – If all packet data has been captured, specific files of interest (for example, file names seen in logs and terminal commands) can potentially be reconstructed from the traffic. These can then be analyzed with specialized tools for the given file format including Portable Executables, PDFs, Java Archives, JavaScript, Microsoft™ OLE documents, and Android™ .apk files.
- **Root Cause Analysis** – When antimalware tools identify and/or analyze files of interest, deep packet data analytics can help with root cause analysis of the file creation. Going beyond just the proximate cause, such as the URL or the dropper that downloaded a particular piece of malware, root cause analysis can determine the entire sequence of events leading to the delivery of the detected malware. This same method of assisted determination can also be performed on data from IDS/IPS, Next-Generation Firewalls (NGFWs), log management platforms and SIEM tools.
- **Validation of Existing Controls** – Analytics capabilities can also be put to good use validating the effectiveness of security controls already in place. For example, packet data can be leveraged to reconstruct sessions in a "before and after" scenario when implementing new firewall or IPS rules.

In a traditional crime, law enforcement gathers existing evidence and then looks for witnesses or security videos of what happened. In an electronic incident, most investigators have very little data that can accurately portray what occurred. Big data analytics provide the definitive proof of what really happened, allowing incident responders to more effectively answer key questions that they may never have been able to answer before.

As with today's new threats, use of new tools and services to accelerate more accurate analysis of large volumes of security, event and deep packet data is also on the rise. However, the question will be, how much data do organizations need to gain context, meet auditor requirements, deliver evidence and give full visibility needed to deal with today's advanced threats? Organizations will need to assure that whatever technological processes they put in place can scale, given that during response, full packets and contents of sessions will need to be reconstructed and even decrypted at the speed investigators need it. Any new layers for analyzing and parsing this traffic will need to work well with existing security, reporting and logging tools.

Conclusion

More and more enterprises and IT security leaders feel they cannot completely prevent compromises and breaches. They are also aware that they may not even know when they are breached, given the subtlety of today's attacks.

While most efforts at prevention reduce incidents, organizations need to be prepared for the inevitability of an attack by having a solid response plan in place for when breaches occur. To do this, IT security and operations professionals need to get better at identifying trends and behaviors that are indicative of normal and abnormal behavior, and then determining the root cause of events in their environment.

As attack trends continue to evolve, the ability to sift through massive amounts of security and event data—while analyzing trends and root cause—will be crucial to risk management, breach preparedness and compliance plans. Without the ability to sort and contextually analyze large quantities of security data, organizations cannot get the visibility into their security and event data they need to determine what was impacted—as well as how it occurred or whether or not the event is really over. These key questions require answers by auditors and responders during and after an event. Answers to these questions are equally important for analyzing and closing the application, network and endpoint vulnerabilities that exposed the organization to the breach and for continuously improving security and operations as new threats dictate.

About the Author

Dave Shackleford, founder and principal consultant with Voodoo Security, is a SANS analyst, instructor and course author, as well as a GIAC technical director. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. He is a VMware vExpert, and has extensive experience designing and configuring secure virtualized infrastructures. He has previously worked as CSO for Configuresoft, CTO for the Center for Internet Security. Dave is the co-author of Hands-On Information Security from Course Technology as well as the "Managing Incident Response" chapter in the Course Technology book *Readings and Cases in the Management of Information Security*. Recently, Dave co-authored the first published course on virtualization security for the SANS Institute. Dave currently serves on the board of directors at the Technology Association of Georgia's Information Security Society and the SANS Technology Institute.

SANS would like to thank its sponsors: